

Self-organized chaos through polyhomeostatic optimization

D. Markovic and Claudio Gros

Institute for Theoretical Physics, Johann Wolfgang Goethe University, Frankfurt am Main, Germany

(Dated: May 31, 2010)

The goal of polyhomeostatic control is to achieve a certain target distribution of behaviors, in contrast to homeostatic regulation which aims at stabilizing a steady-state dynamical state. We consider polyhomeostasis for individual and networks of firing-rate neurons, adapting to achieve target distributions of firing rates maximizing information entropy. We show that any finite polyhomeostatic adaption rate destroys all attractors in Hopfield-like network setups, leading to intermittently bursting behavior and self-organized chaos. The importance of polyhomeostasis to adapting behavior in general is discussed.

Introduction.— Homeostatic regulation plays a central role in all living, as well as in many technical applications. Biological parameters, like the blood sugar level, the heart beating frequency or the average firing rates of neurons need to be maintained within certain ranges in order to guarantee survival. The same holds in the technical regime for the rotation speed of engines and the velocity of airplanes, to give a few examples.

Homeostatic control in the brain goes beyond the regulation of scalar variables like the concentration of proteins and ions, involving the functional stability of neural activity both on the individual as well as on a network level [1–3]. We use here the term ‘polyhomeostasis’ for self-regulating processes aimed at stabilizing a certain target distribution of dynamical behaviors. Polyhomeostasis is an important concept used hitherto mostly implicitly and not yet well studied from the viewpoint of dynamical system theory. Polyhomeostasis is present whenever the goal of the autonomous control is the stabilization of a non-trivial distribution of dynamical states, polyhomeostatic control hence generalizes the concept of homeostasis. The behavior of animals on intermediate time scales, to give an example, may be regarded as polyhomeostatic, aiming at optimizing a distribution of qualitatively different rewards, like food, water and protection; animals are not just trying to maximize a single scalar reward quantity. A concept loosely related to polyhomeostasis is homeokinesis, proposed in the context of closed-loop motion learning [4], having the aim to stabilize non-trivial but steady-state movements of animals and robots.

Here we study generic properties of dynamical systems governed by polyhomeostatic self-regulation using a previously proposed model [5, 6] for regulating the firing-rate distribution of individual neurons based on information-theoretical principles. We show that polyhomeostatic regulation, aiming at stabilizing a specific target distribution of neural activities gives rise to non-trivial dynamical states when recurrent interactions are introduced. We find, in particular, that the introduction of polyhomeostatic control to attractor networks leads to a destruction of all attractors resulting for large networks, as a function of the average firing rate, in either intermittent bursting

behavior or self-organized chaos, with both states being globally attracting in their respective phase spaces.

Firing-rate distributions.— We consider a discrete-time, rate encoding artificial neuron with input $x \in [-\infty, \infty]$, output $y \in [0, 1]$ and a transfer function $g(z)$,

$$y(t+1) = g(a(t)x(t)+b(t)), \quad g(z) = \frac{1}{e^{-z} + 1}. \quad (1)$$

The gain $a(t)$ and the threshold $-b(t)/a(t)$ in (1) are slow variables, their time evolution being determined by polyhomeostatic considerations.

Information is encoded in the brain through the firing states of neurons and it is therefore plausible to postulate [5], that polyhomeostatic adaption for the internal parameters $a(t)$ and $b(t)$ leads to a distribution $p(y)$ for the firing rate striving to encode as much information as possible given the functional form (1) of the transfer function $g(z)$. The normalized exponential distribution

$$p_\lambda(y) = \frac{\lambda e^{-\lambda y}}{1 - e^{-\lambda}}, \quad \mu = \frac{1}{\lambda} \frac{e^\lambda - 1 - \lambda}{e^\lambda - 1}, \quad (2)$$

maximizes the Shannon entropy [7], viz the information content, on the interval $y \in [0, 1]$, for a given expectation value μ . A measure for the closeness of the two probability distributions $p(y)$ and $p_\lambda(y)$ is given by the Kullback-Leibler divergence [7]

$$D_\lambda(a, b) = \int p(y) \log \left(\frac{p(y)}{p_\lambda(y)} \right) dy, \quad (3)$$

which is, through (1), a function of the internal parameters a and b . The Kullback-Leibler divergence is strictly positive and vanishes only when the two distributions are identical. By minimizing $D_\lambda(a, b)$ with respect to a and b one obtains [6] the stochastic gradient rules

$$\begin{aligned} \Delta a &= \epsilon_a \left(1/a + x \Delta \tilde{b} \right) \\ \Delta b &= \epsilon_b \Delta \tilde{b}, \quad \Delta \tilde{b} = 1 - (2 + \lambda)y + \lambda y^2 \end{aligned} \quad (4)$$

which have been called ‘intrinsic plasticities’ [1]. The respective learning rates ϵ_a and ϵ_b are assumed to be small, viz the time evolution of the internal parameters a and b is slow compared to the evolution of both x and y . For

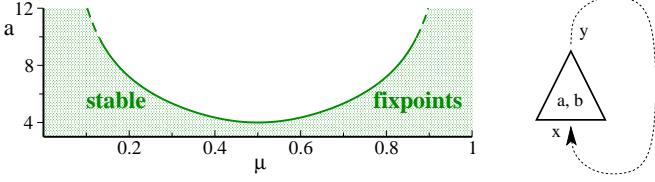


FIG. 1: (color online) Left: As a function of the average firing rate μ , the region of stability (shaded area) of the fixpoint (y^*, b^*) , see Eq. (7) of the one-neuron network in the quasistatic limit. Right: A self-coupling neuron with internal parameters $a(t)$ and $b(t)$.

any externally given time series $x(t)$ for the input, the adaption rules (4) will lead to a distribution of the output firing rates $y(t)$ as close as possible, given the specification (1) for the transfer function, to an exponential with given mean μ .

Attractor relics.— In deriving the stochastic gradient rules (4) it has been assumed that the input $x(t)$ is statistically independent of the output $y(t)$ (but not vice versa). This is not any more the case when a set of polyhomeostatically adapting neurons are mutually interconnected, forming a recurrent network. Here we will show that networks based on polyhomeostatic principles will generically show spontaneous and continuously ongoing activity.

In a first step we analyze systematically the smallest possible network, viz the single-site loop, obtained by feeding the output back to the input, see Fig. 1, a setup which has been termed in a different context ‘autapse’ [8]. We use the balanced substitution $x \rightarrow y - 1/2$ in Eqs. (1) and (4), the complete set of evolution rules for the dynamical variables $y(t)$, $a(t)$ and $b(t)$ is then

$$\begin{aligned} y(t+1) &= g(a(t)[y(t) - 1/2] + b(t)) \\ b(t+1) &= b(t) + \epsilon_b \Delta \tilde{b}(t) \\ a(t+1) &= a(t) + \epsilon_a \left(1/a(t) + [y(t) - 1/2] \Delta \tilde{b}(t) \right) \end{aligned} \quad (5)$$

with

$$\Delta \tilde{b}(t) = 1 - (2 + \lambda)y(t+1) + \lambda y^2(t+1). \quad (6)$$

Note, that $\Delta \tilde{b}(t)$ in (6) depends on $y(t+1)$, and not on $y(t)$, as one can easily verify when going through the derivation of the rules (4) for the intrinsic plasticity. The evolution equations (5) are invariant under $y \leftrightarrow (1 - y)$, $b \leftrightarrow (-b)$, $a \leftrightarrow a$ and $\lambda \leftrightarrow (-\lambda)$, the later corresponding to the interchange of $\mu \leftrightarrow (1 - \mu)$.

We first consider the quasistatic limit $\epsilon_a \ll \epsilon_b$, viz $a(t) \simeq a$ is approximatively constant. The fixpoint (y^*, b^*) in the (y, b) plane is then determined by

$$\begin{aligned} 0 &= \lambda(y^*)^2 - (2 + \lambda)y^* + 1 \\ b^* &= g^{-1}(y^*) - a[y^* - 1/2] \\ &= \log(y^*/(1 - y^*)) - a[y^* - 1/2] \end{aligned} \quad (7)$$

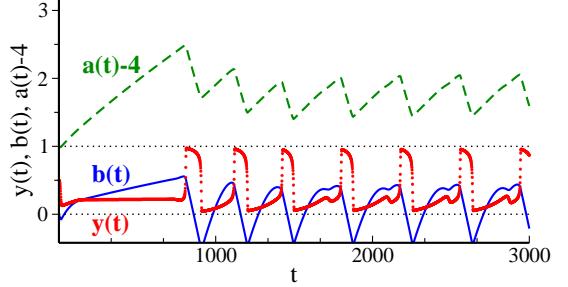


FIG. 2: (color online) The time dependence of $y(t)$ (red stars) $b(t)$ (solid line) and of $a(t) - 4$ (dashed line) for the balanced one-site problem (5). $\epsilon_a = \epsilon_b = \epsilon = 0.01$ and $\mu = 0.28$; the horizontal dotted lines are guides to the eye. The gain $a(t)$ is initially small and the system relaxes, since $\epsilon \ll 1$, fast to a fixpoint of $y = g(a[y - 1/2] + b)$. Once $a(t)$ surpasses a certain threshold, compare Fig. 1, the fixpoint becomes unstable and the system starts to spike spontaneously.

A straightforward linear stability analysis shows, that the fixpoint (y^*, b^*) remains stable for small gains a and becomes unstable for large gains, see Fig. 1. We now go beyond the quasistatic approximation and consider a small but finite polyhomeostatic adaption rate ϵ_a for the gain. Starting with a small gain a we see, compare Eq. (5), that the gain necessarily grows until it hits the boundary towards instability; for small $\Delta \tilde{b}$ the growths of the gain is $a(t) \sim \sqrt{t}$.

In other words, a finite adaption rate ϵ_a for the gain turns the fixpoint attractor (y^*, b^*) into an attractor relic and the resulting dynamics becomes non trivial and autonomously self-sustained. This route towards autonomous neural dynamics is actually an instance of a more general principle. Starting with an attractor dynamical system, viz with a system governed by a finite number of stable fixpoint, one can quite generally turn these fixpoints into unstable attractor ruins by coupling locally to slow degree of freedoms [10, 11], in our case the slow local variables are $a(t)$ and $b(t)$. Interestingly, there are no saddlepoints present in attractor relic networks in general, and in (5) in particular, and hence no unstable heteroclinics, as in a proposed alternative route to transient state dynamics via heteroclinic switching. [12, 13]

In Fig. 2 the time evolution is illustrated for $\lambda = 3.017$, which corresponds to $\mu = 0.28$, see Eq. (2), and $\epsilon_a = \epsilon_b = \epsilon = 0.01$. The system remains in the quasistationary initial regime until the gain a surpasses a certain threshold. The initial quasistationary fixpoint becomes therefore unstable via the same mechanism discussed analytically above, see Eq. (7), for the regime $\epsilon_a \rightarrow 0$, compare Fig. 1. The output activity $y(t)$ oscillates fast between two transient fixpoints of $y = g(a[y - 1/2] + b)$, having a high and a low value respectively. This spiking behavior of the neural activity is driven by spontaneous oscillations in the threshold $-b(t)/a(t)$, shifting the intersection of $g(a[y - 1/2] + b)$ with y forth and back.

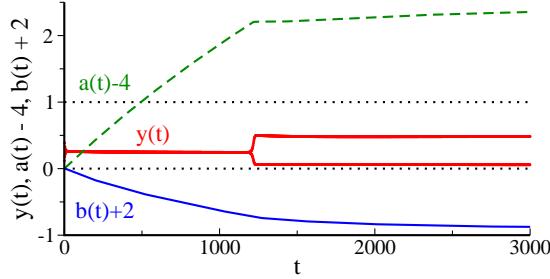


FIG. 3: (color online) The time dependence of $y(t)$ (red stars) $b(t) + 2$ (solid line) and of $a(t) - 4$ (dashed line) for the one-site problem with inhibitory self-coupling $x \rightarrow 1/2 - y$, $\epsilon_a = \epsilon_b = \epsilon = 0.01$ and $\mu = 0.28$. A Hopf-bifurcation occurs for the output $y(t)$ when the initial quasistationary fixpoint becomes unstable, giving place to a new fixpoint of period two.

Evaluating the local Lyapunov exponents we find that the trajectory is stable against perturbations for the transient states close to one of the transient fixpoints of $y = g(a[y - 1/2] + b)$ and sensitive to perturbations and external modulation during the fast transition periods, an instantiation of the general notion of transient state dynamics [10, 11].

Inhibitory self-coupling.— So far we discussed, see (5) and Fig. 2, a neuron having its output $y(t)$ coupled back excitatorily to its input via $x \rightarrow y - 1/2$. The dynamics changes qualitatively for an inhibitory self-coupling $x \rightarrow 1/2 - y$, see Fig. 3. There is now only a single intersection of $g(a[-y + 1/2] + b)$ with y . This intersection corresponds to a stable fixpoint for small gains a . A Hopf-bifurcation [14] occurs when $a(t)$ exceeds a certain threshold and a new fixpoint of period two becomes stable. The coordinates of this fixpoint of period two slowly drift, compare Fig. 3, due to the residual changes in $a(t)$ and $b(t)$. Interestingly, the dynamics remains non-trivial, as a consequence of the continuous adaption, even in the case of inhibitory self-coupling.

Self organized chaos.— We have studied numerically fully connected networks of $i = 1, \dots, N$ polyhomeostatically adapting neurons (4), coupled via

$$x_i(t) = \sum_{j \neq i} w_{ij} y_j(t). \quad (8)$$

The synaptic strengths are $w_{ij} = \pm 1/\sqrt{N-1}$, with inhibition/excitation drawn randomly with equal probability. The adaption rates are $\epsilon_a = \epsilon_b = 0.01$. We consider homogeneous networks where all neurons have the identical μ for the target output distributions (2), with $\mu = 0.15, 0.28$ and $\mu = 0.5$.

The activity patterns presented in the inset of Fig. 4 are chaotic for the $\mu = 0.28$ network and laminar with intermittent bursting for the network with $\mu = 0.15$. This behavior is typical for a wide range of network geometries, distributions of the synaptic weights and for all initial conditions sampled. The respective Kullback-Leibler divergences D_λ decrease with time passing, as shown in

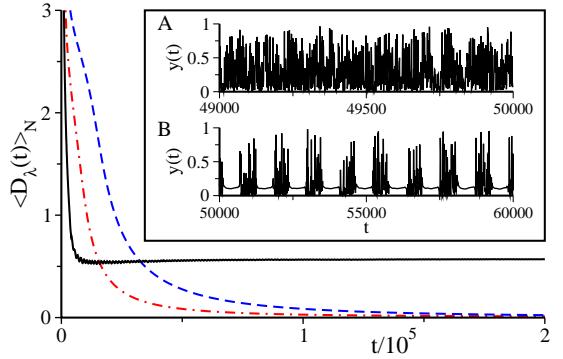


FIG. 4: (color online) Mean Kullback-Leibler divergence D_λ , for $N = 500$ networks, of the real and target output distributions as a function of number of iterations t . The solid black, dashed-dotted red and the dashed blue lines correspond respectively to target firing rates $\mu = 0.15$, $\mu = 0.28$, and $\mu = 0.5$. Inset: Output activity of two randomly chosen neurons for target average firing rates $\mu = 0.28$ (top) and $\mu = 0.15$ (bottom).

Fig. 4, due to the ongoing adaption process (4). D_λ becomes very small, though still finite, for long times in the self-organized chaotic regime ($\mu = 0.28, 0.5$), remaining substantially larger in the intermittent bursting phase ($\mu = 0.15$).

We have evaluated the global Lyapunov exponent, [15] finding that two initial close orbits diverge until their distance in phase space corresponds, within the available phase space, to the distance of two randomly selected states, the tellmark sign of deterministic chaos. [14] The corresponding global Lyapunov exponent is about 5% per time-step for $\mu = 0.5$ and initially increases with the decrease of μ , until the intermittent-bursting regime emerges, after which the global Lyapunov exponent declines with the decrease of μ .

The system enters the chaotic regime, in close analogy to the one-side problem discussed previously, whenever the adaptive dynamics (4) has pushed the individual gains $a_i(t)$, of a sufficient number of neurons, above their respective critical values. Hence, the chaotic state is self-organized. Chaotic dynamics is also observed in the non-adapting limit, with $\epsilon_a, \epsilon_b \rightarrow 0$, whenever the static values of a_i are above the critical value, in agreement with the results of a large- N mean field analysis of an analogous continuous time Hopfield network. [16] Subcritical static a_i lead on the other side to regular dynamics controlled by point attractors.

In Fig. 5 we present the distribution of the output activities for networks with target mean firing rates $\mu = 0.28$ and $\mu = 0.5$. Shown are in both cases the firing rate distributions $p(y)$ of the two neurons having respectively the largest and the smallest Kullback-Leibler divergence (3) with respect to the target exponential distributions (2). Also shown in Fig. 5 are the distributions of the network-averaged output activities.

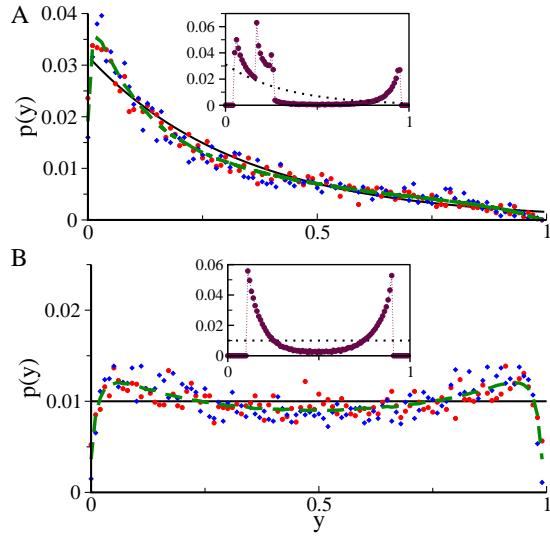


FIG. 5: (color online) Output distributions of the two neurons with highest (blue diamonds) and lowest (red circles) Kullback-Leibler divergence (3) compared to the mean output distribution (dashed green line) and the target exponential output distribution (full black line). For a network with $N = 500$ neurons and a target output mean firing rate (A) $\mu = 0.28$ and (B) $\mu = 0.5$ (identical for all neurons). Insets: For the same parameters the mean output distribution of the single self-coupled neuron, see Fig. 1.

The data presented in Fig. 5 shows, that the polyhomeostatically self-organized state of the network results in firing-rate distributions close to the target distribution. This result is quite remarkable. The polyhomeostatic adaption rules (4) are local, viz every neuron adapts its internal parameters a_i and b_i independently on its own.

Intermittency.— Interestingly, the neural activity presented in the inset of Fig. 4 for $\mu = 0.15$ shows intermittent or bursting behavior. In the quiet laminar periods the neural activity $y(t)$ is close to (but slightly below) the target mean of 0.15, as one would expect for a homeostatically regulated system and the local Lyapunov exponent is negative. The target distribution of firing rates (2) contains however also a small but finite weight for large values for the output $y(t)$, which the system achieves through intermittent bursts of activity. In the laminar/bursting periods the gain $a(t)$ is below/above threshold, acting such as a control parameter [18].

Both the intermittent and the chaotic regime are chaotic in terms of the global Lyapunov being positive in both regimes. A qualitative difference can be observed when considering the subspace of activities (y_1, \dots, y_N) . Evaluating the global Lyapunov exponent in this subspace, viz the relative time evolution $(\Delta y_1, \dots, \Delta y_N)$ of differences in activities, we find this restricted global Lyapunov to be negative in the intermittent and positive in the chaotic regime. Details on the evolution from inter-

mittency to chaos as a function of μ and system sizes will be given elsewhere.

Conclusions.— It is well known [17], that individual and networks of spiking neurons may show bursting behavior. Here we showed, that networks of rate encoding neurons are intermittently bursting for low average firing rates, when polyhomeostatically adapting, entering a fully chaotic regime for larger average activity level. Autonomous and polyhomeostatic self-control of dynamical systems may hence lead quite in general to non-trivial dynamical states and novel phenomena.

Polyhomeostatic optimization is of relevance in a range of fields. Here we have discussed it in the framework of dynamical system theory. Alternative examples are resource allocation problems, for which a given limited resource, like time, needs to be allocated to a set of uses, with the target distribution being the percentages of resource allocated to the respective uses. These ramifications of polyhomeostatic optimization will be discussed in further studies.

References

- [1] G.G. Turrigiano, Trends Neurosci **22**, 221 (1999).
- [2] G.W. Davis, Annual review of neuroscience **29**, 307 (2006).
- [3] E. Marder, and J.M. Goaillard, Nature Reviews Neuroscience **7**, 563 (2006).
- [4] R. Der, U. Steinmetz, and F. Pasemann, in *Computational Intelligence for Modelling, Control, and Automation, Concurrent Systems Engineering* **55**, 43 (1999).
- [5] M. Stemmler, C. Koch, Nature Neuroscience **2**, 521 (1999).
- [6] J. Triesch, in *Proceedings of ICANN 2005*, W. Duch et al. (Eds.), LNCS **3696**, 65 (2005).
- [7] T.M. Cover, J.A. Thomas, *Elements of information theory*, Wiley 2006.
- [8] H.S. Seung, D.D. Lee, B.Y. Reis, and D.W. Tank, J. Comp. Neuroscience **9**, 171 (2000).
- [9] J.H. Steil, Neural Networks **20**, 353 (2007).
- [10] C. Gros, New Journal of Physics **9**, 109 (2007).
- [11] C. Gros, Cognitive Computation **1**, 77 (2009).
- [12] M. Rabinovich *et al*, Phys. Rev. Lett. **87**, 68102 (2001).
- [13] C. Kirst, and M. Timme, Phys. Rev. E **78**, 065201 (2008).
- [14] C. Gros, *Complex and Adaptive Dynamical Systems, A Primer*, Springer (2008).
- [15] R. Ding, and J. Li, Phys. Lett. A **364**, 396 (2007).
- [16] H. Sompolinsky, A. Crisanti, and H.H. Sommers, Phys. Rev. Lett. **61** 259 (1988).
- [17] M.A. Arbib, *The handbook of brain theory and neural networks*, MIT Press (2002).
- [18] N. Platt, E.A. Spiegel, and C. Tresser, Phys. Rev. Lett. **70**, 279 (1993).